

Robotics Research Technical Report

Generatorium omnis labors ex machina

Solving Ill-Conditioned Problems by Minimizing Equation Error

by

Robert A. Hummel †

Courant Institute of Mathematical Sciences

Robert Moniot ‡

Fordham University

Technical Report No. 267


Robotics Report No. 93

December, 1986

NYU COMPSCI TR-267 c.1
Hummel, Robert A.
Solving ill-conditioned
problems by minimizing
equation error.

Courant Institute of Mathematical Sciences
New York University

Computer Science Division
Mercer Street New York, N.Y. 10012



Solving Ill-Conditioned Problems by Minimizing Equation Error

by

Robert A. Hummel †
Courant Institute of Mathematical Sciences

Robert Moniot ‡
Fordham University

Technical Report No. 267
Robotics Report No. 93
December, 1986

†New York University
251 Mercer Street
New York, NY 10012
hummel@nyu.arpa

‡Division of Science and Mathematics
Fordham University
College at Lincoln Center
New York, NY 10023
moniot@nyu-acf8.arpa

Submitted to the First International Conference on Computer Vision
to be held June, 1987

This research was supported by Office of Naval Research Grant N00014-85-K-0077, Work Unit NR 4007006, and NSF grant DCR-8403300. We thank Robert Kohn for providing many of the ideas, and for his helpful guidance.

Solving Ill-Conditioned Problems by Minimizing Equation Error

Robert Hummel
Robert Moniot

Abstract

Ill-conditioned problems arise frequently in computer vision because the image information contains noise and ambiguities such that the true identity of the scene is not uniquely specified. It is nonetheless possible for humans to infer the structure of objects from incomplete data, presumably through the use of assumptions and computations consistent with natural images. We would like to understand these assumptions and computations in order to develop robust computer vision algorithms.

This paper focuses on a method in numerical analysis for solving “inverse problems” when the degraded scene has undergone a sequence of steps modeled by a known equation. Reconstruction can then be attempted by finding a solution that minimizes the deviation from that equation. The method is exemplified by its application to the deblurring problem. In this problem, deblurring is achieved by computing a succession of images, each slightly deblurred from the previous, such that the complete set satisfies the equations specifying the diffusion process of blurring. The method can be viewed as a new approach to regularization for problems where a scale-space parameter can be used to separate the information extracted from the image.

1. Overview

Many problems in computer vision involve the analysis of physical objects and surfaces given relatively little sensed information about those objects. A complete representation of objects implies an ability to reconstruct (or predict) a scene. Therefore, it is important to be able to infer much of the missing information through the use of assumptions about image data. Poggio and Torre [1] argue that much of computational vision is concerned with regularization methods applied to ill-posed problems arising in vision due to insufficient numbers of constraints derived from the visually sensed data. Regularization methods form a class of techniques for attacking optimization and other variational problems, frequently encountered in computer vision. Properly applied, an ill-posed problem becomes well-posed, although sometimes the problem is transformed such that a solution to the regularized problem is no longer a solution to the ill-posed problem. With

regularization methods, generally some form of smoothness assumption is incorporated explicitly into the solution process.

Images, however, are not necessarily smooth. They can exhibit discontinuous behavior, both in terms of intensity and other features, at different spatial scales and at different quantization scales. In some cases, a model of an image as a finite collection of smooth regions separated by discontinuities is excessively constraining, and forbids the extraction or analysis of ranges of textures, edges, and other primitives confounded simultaneously at multiple scales and at common locations.

In David Marr's computational theory of vision, the "second physical assumption" concerns the organization of information at different scales of spatial resolution [2]. In the extraction of zero-crossings, grouping of similar features in the primal sketch, and in algorithms using the primal sketch, the scale of resolution is explicit in the image computation. For example, zero-crossings are computed using a Laplacian-of-Gaussian filter against the image data; the spread of the Gaussian encodes the scale of the feature extraction. Although it is clear that the classification of primitives by scale is useful, it is less clear how to combine information derived at different scales. Marr suggests, for example, that zero-crossing segments that persist at the same location through several scales are associated with physical edges. In the use of the primal sketch for stereo analysis, information from disparities at coarse resolution levels is used to bias the search at fine resolutions, invoking a coarse-to-fine information flow. More generally, however, we would like to use information from all scales without imposing a direction of deduction.

A robust means of inferring structure in images from sparse or underconstrained information, possibly supplied at varying scales, will facilitate a number of useful tasks. In this paper, we concentrate on the problem of deblurring blurred image data. However, this same task can be regarded as a feature extraction process, whereby discontinuities and textures are amplified by deblurring to make their structure more explicit. Indeed, deblurring a normal image can provide useful information for a segmentation of the scene. Other vision problems can be attacked directly by the same methods. For instance, interpolation of sparse data can be handled as a scale-space reconstruction problem, where the sparse data supplies constraints at the appropriate (usually fine-resolution) scale. The deblurring problem, as posed later in this paper, is identical, except that constraints are supplied at one coarse resolution level. Sometimes the data to be interpolated includes gradient information, and this data should also be includable in the reconstruction problem.

Another application area concerns the investigation of the information content of zero-crossings in the raw primal sketch. Marr speculated on the possibility of reconstructing image data given just the zero-crossing information in scale space [2], and others have since shown some favorable examples, doing reconstructions from zero-crossing at a single scale [3]. It is known that reconstruction is theoretically possible, given zero-crossings and gradient information along the zero-crossings [4], but good, stable, reconstruction methods are mostly unexplored. The interest in reconstructions, of course, is mostly related to a desire to know what additional assumptions about images makes the representation stable. This, in turn, can guide the further extraction of salient information from the representation, such as in the grouping of zero-crossings into recognized structures.

In this paper, we outline a new method for tackling these kinds of problems, based on a method for the solution of ill-conditioned problems. Different methods will yield different solutions, and the “correct” solution depends on the nature of human perception and expectations about the problem domain. The approach taken here, minimization of equation error, seems to have an intuitive appeal, and based on a few experiments, great promise. A fundamental tenet of the approach is that information and assumptions about images are injected in a graded family of multiresolution representations of the data. Accordingly, the range of applications will be limited to problems that have natural multiresolution formulations.

In the next section, we discuss various methods for solving the ill-conditioned deblurring problem, in order to place minimization of equation error (Method (5)) among the alternatives. Section 3 gives some mathematical details of the approach, and Section 4 gives results of some preliminary experiments.

2. Methods for solving ill-conditioned problems

Consider the problem of deblurring a blurred signal. That is, we wish to determine f given $b = G * f$, where G is a Gaussian convolution kernel. The problem is well-defined, given certain restrictions on b and the space of admissible functions for f . This is because distinct f 's give rise to distinct b 's. However, the problem is still ill-posed, due to the lack of stability. The difficulty is that widely different signals, f_1 and f_2 , can yield blurred signals b_1 and b_2 that are arbitrarily close together. Thus minor inaccuracies in the representation of b can lead to large inaccuracies in the deblurred function f .

However, if we restrict f to the set of naturally-occurring signals, then the situation is not necessarily hopeless. In many ill-posed vision problems, including deblurring, it may be possible to make certain assumptions about natural images such that reconstruction is in fact possible. The advantage of studying reconstructions is that an understanding of the assumptions underlying natural images might transfer to methods for interpreting and inferring structure in sensed image data.

Various methods have been studied for solving deblurring problems with natural images. There exist several classes of methods for regularizing the deblurring problem. Specifically, we are given b , and we seek f satisfying $G * f = b$. We have the following alternatives:

- (1) We can restrict the space of admissible f 's, and seek such an f minimizing

$$\|G * f - b\|^2.$$

- (2) We define some process $\mathcal{P}b$ that computes a solution to the problem $G * f = b$ among a restricted class of functions f . We then apply the process to any given b , even though the initial data f may not satisfy the restrictions. The resulting $\mathcal{P}b$ may not be an exact deblurring, but may be good enough.
- (3) We discretize the problem so that many f 's satisfy $G * f = b$ to machine precision. We then seek a solution f among the set of all solutions minimizing

$$\|f\|^2,$$

where some norm $\|\cdot\|$ has been defined on the space of admissible functions. This is the least squares solution.

- (4) We can use a regularization term, to seek an f minimizing

$$\|G * f - b\|^2 + \epsilon \|f\|^2.$$

The solution, f_ϵ , will be an approximate solution to the deblurring problem for small ϵ . Depending on the norm $\|\cdot\|$, it can sometimes be shown that f_ϵ converges to a solution f as $\epsilon \rightarrow 0$.

Minimization of equation error would prescribe the following method:

- (5) Observe that the solution to the initial value Heat Equation,

$$\Delta u = \frac{\partial u}{\partial t}$$

$$u(x,0) = f(x),$$

yields the blurred signal b at some scale $u(x,T)=b(x)$. Thus, find a solution $u(x,t)$ satisfying $u(x,T)=b(x)$ minimizing error in the equation

$$\|\Delta u - \frac{\partial u}{\partial t}\|^2.$$

Note that the norm is measured over all (x,t) space. The deblurred signal is extracted as $f(x)=u(x,0)$.

Deblurring is the most extensively studied ill-conditioned problem. Method (1) appears in many forms. John [5] considered functions f satisfying $f \geq 0$; a solution among the class of band-limited functions with fixed maximum spectral frequency is easy; solutions among spaces of splines or polynomials are also common. Shvaytser and Peleg [6] study a conjugate gradient minimization approach for the case $0 \leq f \leq M$. Carasso *et al* [7] solve deblurring among f satisfying $\|f\|_{L^2}^2 \leq M$. Method (2) is the basis of [8], where the process used is the one that precisely deblurs polynomials of small degree. Method (3) can also lead to spline solutions, and in general makes explicit smoothness assumptions. Method (4) has been studied extensively in many domains; it is the basis of the continuity method for the solution of non-viscous fluid flow problems [9]; Poggio and others have used similar methods for studying smoothness constraints [10].

Our interest in the last method, method (5), is inspired by results of Bruce Lowe in a recent thesis [11]. His thesis advisor, Robert Kohn, explains that the variational approach to certain elliptic identification problems seems to yield computational advantages and far superior results over common numerical procedures. Despite its simplicity, the approach of minimizing equation error is relatively new, and in Lowe's thesis, is applied to problems of steady-state porous media flow (oil flow in subterranean strata). Kohn suggested that these methods might be applicable to parabolic problems, and our experiments were based on this suggestion.

One possible reason why these methods are so new, despite their apparent likelihood for success, is that the memory requirements of the approach are very large. In order to minimize equation error, a separate image is needed for each intermediate stage of deblurring; memory requirements of this magnitude were infeasible very few years ago.

Why should we prefer Method (5) over the others? To date, our interest is founded purely on empirical grounds. Since the topic is related to natural images,

evaluation of the results will necessarily rely on a large subjective component. There is a separate issue as to the mathematical explanations of the method's success and weaknesses. We suspect that this analysis will be difficult and needs to be guided by empirical results. It is easy to prove, mathematically, that the minimization process will converge, and that if a solution exists, equation error minimization will find a solution. However, it is a much different thing to show that the process finds a "natural" or "correct" solution when the data is from a natural scene. Here, we justify our preference for the proposed method on a few results, and with an appeal to certain intuitive attractions of the method of minimizing equation error.

Specifically, there are no explicit smoothness assumptions, so that the resulting deblurred signal can have sharp discontinuities. The deblurring takes place incrementally, through a sequence of stages, although the flow of information is not constrained to follow from blurred-to-fine resolution scales. In essence, the intermediate levels of deblurring impose a degree of regularization on the process. However, the minimization of equation error is a fundamentally different process from other regularization methods. For example, the term $\|G*f-b\|^2$ does not appear in the equation error. A solution u will necessarily satisfy $u(\cdot, T)=b$, but does not necessarily satisfy $u(\cdot, T)=u(\cdot, 0)*G$ exactly; instead the error is spread throughout all scales. A final attraction of the method involving multiple scales of blurring is that information (constraints) can be injected into the system at different scales. In the deblurring problem, the data is specified at the level $t=T$; however, there exist many vision problems where sparse information is supplied at different levels.

One way to view the minimization of equation error as opposed to more traditional regularization methods is to consider the parabolic Heat Equation as a dynamical system. Then if the data at any fixed level $u(\cdot, t)$ is thought of as a vector function evaluated at t , $\bar{u}(t)$, and the Laplacian operator is a vector-valued transformation on the vector data $f(\bar{v})$, the Heat Equation becomes

$$\frac{d\bar{u}}{dt} = \bar{f}(\bar{u}) .$$

Minimizing equation error is akin to a standard numerical relaxation procedure for solving ordinary differential equations, minimizing

$$\int \left\| \frac{d\bar{u}}{dt} - \bar{f}(\bar{u}) \right\|^2 dt .$$

Regularization methods, on the other hand, are similar to "shooting methods," where the boundary conditions at one end are determined so that the resulting trajectory satisfies the boundary conditions at the other end.

3. Minimizing Equation Error

In this section, we give some of the details in the formulation of the minimization of equation error, and discuss a few details of implementation. The prototypic ill-conditioned problem is the task of deblurring a blurred signal. Accordingly, we describe the idea of minimizing equation error in the context of the deblurring problem.

3.1. Data Error and Equation Error

We denote the initial signal, with domain in \mathbb{R}^n , by f , and the blurred signal by $\mathcal{T}f$. In the unbounded domain \mathbb{R}^n , $\mathcal{T}f$ is given by Gaussian blurring of f :

$$(\mathcal{T}f)(x) = \frac{1}{(4\pi T)^{n/2}} \int f(x') e^{-\frac{|x-x'|^2}{4T}} dx' ,$$

for some fixed constant $T > 0$. It is an essential component to the equation error approach that there exists a graded set of intermediate blurrings $u(x, t)$, parameterized by t , $0 \leq t \leq T$,

$$u(x, t) = \frac{1}{(4\pi t)^{n/2}} \int f(x') e^{-\frac{|x-x'|^2}{4t}} dx' .$$

The parameterized blurrings $u(x, t)$ satisfy the initial value Heat Equation:

$$\Delta u = \frac{\partial u}{\partial t} , \quad t > 0 , \quad x \in \mathbb{R}^n .$$

$$u(x, 0) = f(x) , \quad x \in \mathbb{R}^n .$$

Similar formulations can be given when the domain of the signal $f(x)$ is bounded, and boundary conditions must be imposed. In all cases, however, the Heat Equation dictates the blurring process in the parameterized scale-space $0 \leq t \leq T$.

The deblurring problem is given as follows. We are given a blurred signal $b(x)$, and perhaps an estimate of the extent of blur T . We seek a deblurred signal $f(x)$ such that

$$\mathcal{T}f = b .$$

One way of posing this problem, in practice, is to define a normed space of admissible functions $f \in \mathcal{F}$, and seek

$$f \in \mathcal{F} \text{ minimizing } \|\mathcal{T}f - b\| .$$

We refer to this formulation as the problem of minimizing *data error*, since b is the given data.

An alternate formulation, to be studied here, focuses on the graded set of blurrings $u(x, t)$. We again pose a normed space of admissible functions $u \in \mathcal{E}$, which include (as part of the definition of \mathcal{E}) the requirement that

$$u(x, T) = b(x) , \quad \text{for } u \in \mathcal{E} .$$

Note that the domain of the functions in \mathcal{E} is the entire scale-space (x, t) , $x \in \mathbb{R}^n$ and $0 \leq t \leq T$. Noting that u should satisfy the Heat Equation we solve the problem

$$\text{Find } u \in \mathcal{E} \text{ minimizing } \left\| \Delta u - \frac{\partial u}{\partial t} \right\| .$$

This is an entirely different formulation. Minimizing equation error involves minimizing an integral over scale-space; e.g.,

$$\|\Delta u - \frac{\partial u}{\partial t}\|^2 = \int_0^T \int_{\mathbb{R}^2} |\Delta u - \frac{\partial u}{\partial t}|^2 dx dt .$$

Minimizing data error is simply an integral over \mathbb{R}^n , such as

$$\|\mathcal{T}f - b\|^2 = \int_{\mathbb{R}^n} |\mathcal{T}f(x) - b(x)|^2 dx .$$

The former is not simply an extension of the latter, since the data error term does not appear in the equation error term.

We in fact advocate a slightly different measure of the equation error. We transform the Heat Equation, a second order partial differential equation, into a first order system:

$$\nabla u = \bar{\sigma} ,$$

$$\nabla \cdot \bar{\sigma} = \frac{\partial u}{\partial t} .$$

We define a new function space \mathcal{G} consisting of tuples of functions $(u, \bar{\sigma}) \in \mathcal{G}$ with a norm of the form

$$\|(u, \bar{\sigma})\| = \int_0^T |u|^2 + |\bar{\sigma}|^2 dx dt .$$

Included in the definition of \mathcal{G} are the boundary conditions

$$u(x, T) = b(x) .$$

The equation error of an element $(u, \bar{\sigma}) \in \mathcal{G}$ is the norm

$$\|(\nabla \cdot \bar{\sigma} - \frac{\partial u}{\partial t}, \nabla u - \bar{\sigma})\| .$$

We then seek an element $(u, \bar{\sigma}) \in \mathcal{G}$ minimizing the equation error. That is, we find $(u, \bar{\sigma})$ with $u(x, T) = b(x)$ minimizing

$$\int_0^T \int_{\mathbb{R}^n} |\nabla \cdot \bar{\sigma} - \frac{\partial u}{\partial t}|^2 + |\nabla u - \bar{\sigma}|^2 dx dt . \quad (2.1)$$

3.2. Discrete Formulation

A discretization of these ideas can be performed in a number of different ways. We will describe, briefly, a finite element and a finite difference approach. Most of our discussion is limited to one spatial dimension ($n=1$), in which case $\bar{\sigma}$ is a scalar function.

Scale-space (the domain of the Heat Equation) has variables (x, t) , $x \in \mathbb{R}$, $t \geq 0$, when $n=1$. We assume a regular rectangular discrete grid, $(i \cdot h, k \cdot s)$, $i = \dots, -1, 0, 1, \dots$, and $k = 0, 1, \dots$. For a finite element formulation [12], the simplest elements can be formed by splitting each grid square into two triangular regions, along one of the two chosen diagonals (standard triangular elements); see Figure 1. The unknown functions, u and σ can be represented as continuous piecewise linear functions, linear within each triangular region, by simply specifying the value of the functions at each grid point. Accordingly, the unknowns are the values




Figure 1. Triangular, piecewise-linear, finite elements.

of $u(ih, ks)$ and $\sigma(ih, ks)$, for all i and k , and will be denoted by $u_{i,k}$ and $\sigma_{i,k}$. The simple piecewise linear elements suffice since the equation was converted to a first order system. It is then a simple (albeit tedious) matter to compute the Equation Error (Equation 2.1) for such functions in terms of the unknowns. It turns out that the equation error is a quadratic function in the unknowns, and that there is locality of dependence. That is, the gradient of the equation error depends linearly (via a matrix) on the unknowns, and that a component of the gradient with respect to an unknown at (i, k) depends only on the values at nearby grid points.

The case $h=1/2, s=1$ is very common, and so we will give the gradient equations for this case. The result can be represented as a sum of discrete convolutions, with the data u and σ considered as image arrays. Then for E equal to the equation error, the values of $\partial E / \partial u_{i,k}$ and $\partial E / \partial \sigma_{i,k}$ form image arrays also, which we will denote by $\nabla_u E$ and $\nabla_\sigma E$. The formulas (valid for $k > 0$) are

$$\nabla_u E = \begin{bmatrix} 0 & -4 & 0 \\ -1 & 10 & -1 \\ 0 & -4 & 0 \end{bmatrix} * u + \frac{1}{3} \begin{bmatrix} -4 & 4 & 0 \\ 1 & -6 & 5 \\ 0 & 2 & -2 \end{bmatrix} * \sigma,$$

and

$$\nabla_{\sigma} E = \frac{1}{3} \begin{bmatrix} -2 & 2 & 0 \\ 5 & -6 & 1 \\ 0 & 4 & -4 \end{bmatrix} * u + \frac{1}{3} \begin{bmatrix} 1 & 1 & 0 \\ -2 & 12 & -2 \\ 0 & 1 & 1 \end{bmatrix} * \sigma .$$

Here, the three-by-three arrays are masks that are employed in three-by-three local convolutions against the array data u and σ . Separate equations are needed for $k=0$ ($\partial E/\partial u_{i0}$ and $\partial E/\partial \sigma_{i0}$) and on the borders, if boundary conditions are imposed for $i=\pm M$. Note that a certain amount of asymmetry results simply from choosing triangular elements based on one of two possible diagonal cuts through each rectangular grid.

A finite difference approach can also be taken. In this formulation, the domain is again discretized on a rectangular grid, and the partial derivative is approximated by differences of function values at the grid points.

The Heat Equation can be discretized by the approximations

$$(\Delta u)_{i,k} = \frac{1}{4} (u_{i-1,k} - 2u_{i,k} + u_{i+1,k})$$

$$\left(\frac{\partial u}{\partial t} \right)_{i,k} = u_{i,k+1} - u_{i,k} ,$$

yielding the discrete Heat Equation

$$u_{i,k+1} = \frac{1}{4}u_{i-1,k} + \frac{1}{2}u_{i,k} + \frac{1}{4}u_{i+1,k} .$$

This is the numerical method used to compute the blurred signal. Note that, implicitly, we have set $h=1/2$ and $s=1$. However, we once again favor a first order system in place of the second order Heat Equation, and so we introduce variables σ_{ik} to represent values for $\partial u/\partial x$ evaluated at the locations $((i-1/2)h, ks)$. Note that the σ -variables are offset a half-pixel in the horizontal direction, so that the equation $\nabla u = \sigma$ becomes

$$\frac{u_{i,k} - u_{i-1,k}}{2} = \sigma_{i,k} .$$

The other equation, $\nabla \cdot \sigma = \partial u/\partial t$, becomes

$$u_{i,k+1} - u_{i,k} = \frac{\sigma_{i+1,k} - \sigma_{i,k}}{2} .$$

Thus the equation error is

$$E = \sum_{k \geq 0} \sum_i \left(\frac{u_{i,k} - u_{i-1,k}}{2} - \sigma_{i,k} \right)^2 + \left(u_{i,k+1} - u_{i,k} - \frac{\sigma_{i+1,k} - \sigma_{i,k}}{2} \right)^2 .$$

For u a solution to the discrete Heat Equation, the total equation error can be made zero by a suitable choice of the σ -variables.

Clearly, the equation error for the finite difference formulation is once again a quadratic function of the variables. Minimization of equation error will involve the gradient of E with respect to the u and σ variables. Once again, except near border pixels, these gradient values, $\partial E/\partial u_{i,k}$ and $\partial E/\partial \sigma_{i,k}$ can be viewed as images obtained by convolutions against the grid arrays $u_{i,k}$ and $\sigma_{i,k}$. The formulas are:

$$\nabla_u E = \begin{bmatrix} 0 & -4 & 0 \\ -1 & 10 & -1 \\ 0 & -4 & 0 \end{bmatrix} * u + \begin{bmatrix} 0 & 0 & 0 \\ 0 & -4 & 4 \\ 0 & 2 & -2 \end{bmatrix} * \sigma ,$$

$$\nabla_\sigma E = \begin{bmatrix} -2 & 2 & 0 \\ 4 & -4 & 0 \\ 0 & 0 & 0 \end{bmatrix} * u + \begin{bmatrix} 0 & 0 & 0 \\ -1 & 6 & -1 \\ 0 & 0 & 0 \end{bmatrix} * \sigma .$$

The finite element and finite difference formulations lead to slightly different quadratic functionals for the equation error. We suspect they will behave similarly, although this is one topic for further investigation.

3.3. Boundary Conditions

Since images are typically defined on a bounded domain, we must consider how to handle border conditions. We will discuss three possibilities, and describe the impact on the discrete formulations given in the previous subsection.

A simple method for handling borders is to extend the image by zeros outside its original domain, and allow the support of the blurred image to expand by one pixel in each blurring step. As long as the extension is sufficiently large relative to the number of blurring steps, the borders can be ignored by this method. The u and σ variables can be pegged at zero in the appropriate extended regions.

We can also impose boundary conditions motivated by the analogy to the Heat Equation. The two natural choices are Dirichlet boundary conditions (on the spatial borders) and Neumann conditions. In terms of the Heat Equation defined in the scale-space $\{(x, t) | x \in \mathcal{D}, t \geq 0\}$, where \mathcal{D} is the image domain, Dirichlet boundary conditions amount to specifying

$$u(x, t) = f(x), \quad x \in \partial\mathcal{D}, \quad t \geq 0 ,$$

whereas Neumann conditions are given by

$$\frac{\partial u}{\partial \nu}(x, t) = 0, \quad x \in \partial\mathcal{D}, \quad t \geq 0 ,$$

where $\partial u / \partial \nu$ is outward normal spatial derivative. In both cases, other constraints may be imposed, and are unaffected by these border considerations. For example, in the deblurring problem, we have the constraint $u(x, T) = b(x)$, $x \in \mathcal{D}$. Note that for Dirichlet conditions, however, it is assumed that the border values of the image data, $f(x)$, $x \in \partial\mathcal{D}$, are known. Incidentally, the Dirichlet border conditions correspond physically to heat diffusion of a system in contact with fixed-temperature heat reservoirs, whereas Neumann conditions arise in adiabatic (insulated) diffusion.

The discrete formulation of equation error can be modified to account for either of these border conditions. For both the borders $x \in \partial\mathcal{D}$ and the bottom border $t=0$, in both boundary formulations, extra data points are added in a single layer outside the border. The u and σ variables at these extra points are not free, but determined according to rules that depend on the boundary condition. For example, for the Neumann conditions on the side borders, we require the values of u along the added border points to equal their neighbors in the interior, and the σ

values located between those grid points are set to zero. Along the bottom $t=0$, we similarly add a row of data points for u (corresponding to $t=-1$), and a row of σ values. The u -row is filled with a copy of the $t=0$ data, and the added σ -row is held fixed with zeros.

By adding extra data points in this way, it turns out that the equations for the gradient of the Equation Error can be applied precisely as stated before, applicable to all pixels except the added points, without additional terms.

3.4. Two dimensions

The case $n=2$, i.e., the image domain \mathcal{D} is a subset of \mathbb{R}^2 , can be handled by similar methods, but requires additional design choices. A fairly natural discrete finite difference approach proceeds with introduction of a grid of points (i, j, k) , $k \geq 0$, and variables $u_{i,j,k}$, $\sigma_{i,j,k}^{(1)}$, $\sigma_{i,j,k}^{(2)}$, and $v_{i,j,k}$, with the system of equations

$$\begin{aligned}\sigma_{i,j,k}^{(1)} &= \frac{u_{i,j,k} - u_{i-1,j,k}}{2}, \\ \sigma_{i,j,k}^{(2)} &= \frac{u_{i,j,k} - u_{i,j-1,k}}{2}, \\ v_{i,j,k} - u_{i,j,k} &= \frac{\sigma_{i+1,j,k}^{(1)} - \sigma_{i,j,k}^{(1)}}{2},\end{aligned}$$

and

$$u_{i,j,k+1} - v_{i,j,k} = \frac{\sigma_{i,j+1,k}^{(2)} - \sigma_{i,j,k}^{(2)}}{2}.$$

When these equations hold exactly, the data u is blurred from level k to level $k+1$ by convolution against the kernel:

$$\frac{1}{16} \begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix}.$$

More generally, the four equations above can be converted to a quadratic measure of equation error in the unknown variables. Gradient equations can be formed, and will once again be linear with local (three-by-three-by-three) dependence. The gradient equations can be used in either a steepest descent or in a conjugate gradient procedure to minimize equation error.

3.5. Minimization Techniques

The equation error E for the discretized deblurring problem can be minimized by any of a number of standard techniques. The situation is especially simple since E is a quadratic functional. Since the gradient is easily computed, our initial experiments were with the method of steepest descent. As an initial estimate for this iterative approach, we simply filled all of the scale-space with the blurred data. Jacobi overrelaxation was used, and satisfactory results have been observed with small data sets (one spatial dimension, a dozen discretization points in that dimension, and a half dozen blurring steps). However, as one might expect, convergence is extremely slow with even moderate size data sets.

We have thus also tried conjugate-gradient methods for minimization of equation error. This method has the drawback of requiring additional storage space, increases computational requirements per iteration, and also requires a term which constitutes a global sum over all grid points. However, because E is quadratic, the method has (theoretically) finite time convergence, requiring no more iterations than the number of unknowns. In practice, roundoff error, even with double-precision calculations, necessitates additional iterations for exact convergence. Nonetheless, the rate of convergence for the conjugate-gradient method is vastly superior to the steepest descent method.

We have not yet tried the “pre-conditioned conjugate-gradient” method for minimizing the quadratic equation error [12-14]. We would expect this method to greatly speed the computation of a minimum in equation error. However, due to the ill-conditioned nature of the problem, pre-conditioning might change the character of the solution, so it is not certain that the method will still yield good results.

4. Some Experimental Results

After some simple experiments using small test data sets, we conducted a number of runs using 512 data points taken from a scanline of an image, and some simple examples using 64 by 64 pixel images.

Figure 2 shows plots of original data (an image scanline), the same data blurred by 10 steps of blurring using a $[1/4, 1/2, 1/4]$ blurring kernel and Neumann boundary conditions (this corresponds approximately to Gaussian blur with standard deviation equal to 2.2 pixels), and the results of deblurring by minimizing equation error. For the deblurring procedure, the finite difference formulation of equation error was used, and conjugate gradient minimization methods were necessary. There are approximately 10,000 degrees of freedom in the system, and one expects exact convergence after 10,000 iterations. In fact, Figure 2 represents roughly 20,000 iterations, although the amount of change after 1000 iterations is very slight. Note that the deblurred signal accurately reconstructs many of the features of the original signal.

We have tried the same process using other signals as input data, with similar results. We also programmed the case for Dirichlet boundary conditions, and achieve similar deblurrings of data blurred using Dirichlet boundary conditions (i.e., with the data on the left and right edge fixed). Finally, we applied the deblurring process formulated using Neumann boundary conditions to data that had been blurred using Dirichlet boundary data. The result was that satisfactory deblurring was obtained, although more experiments are needed to say anything definitive.

All calculations were done using double-precision floating point calculations. To analyze the extent to which the success depends on the high degree of accuracy in the representation of the data, we repeated the deblurring experiments with noise added to the blurred data. Random noise of maximum amplitude 10^{-5} was added to the blurred data. Here, the original data lies in the range 0 to 1. Thus this experiment is in part an indication of the best that can be expected with single-precision floating point arithmetic. The results are still excellent; there was hardly any difference from Figure 2.




Figure 2. An example of deblurring using minimization of equation error.




Figure 3. The results of deblurring when noise of magnitude $1/256$ has been added to the blurred signal.




Figure 4. An original 64 by 64 image (upper left), the same image blurred using a Gaussian with standard deviation equal to two pixels (upper right), and two stages of deblurring, one at an intermediate level in scale space (bottom left), and the other at the base (deblurred) level (bottom right).

We tried the same experiment, adding noise of magnitude $1/256$. Our interest here is in the effects of representing the blurred data with single bytes. The results are shown in Figure 3. We observe considerable degradation in the performance of the deblurring, but were nonetheless encouraged by the robustness of the change.

Finally, we attempted to deblur image data. The image we used is of size 64 by 64 pixels, subsampled from a portion of a larger image. Eight stages of blurring using the kernel

$$\frac{1}{16} \begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix}$$

were applied (corresponding to Gaussian blur with standard deviation 2.0 pixels). Neumann boundary conditions were used. A two-dimensional version of the finite-difference form of the equation error was formulated, and programs to perform conjugate gradient minimization were written. We chose a formulation slightly different from the one described in Section 3.4, although it is not likely that the precise formulation will have much influence. We also used the layering technique described below to reduce the computational expense entailed. The results are depicted in Figure 4. (True grayscale images, as opposed to the laser-printed simulated halftones given here, are slightly more revealing.) We show the original, blurred, partly deblurred, and fully deblurred images. The partly deblurred image was taken from level four of the eight levels of the scale-space data.

We also tried deblurring an image that was blurred by optical means¹ instead of by numerical blurring. Figure 5 shows the blurred and deblurred images.

The deblurring of the 64 by 64 images took an enormous amount of computer time (hours). We therefore experimented with a layering technique, where we first deblur half of the layers, and then use the partly deblurred data as the fixed data for the second half of the layers. Each of the two slabs requires (at least theoretically) one-fourth of the amount of work of the single large slab (using a sequential processor model of computation), and roughly half of the memory. Thus the total amount of work required is cut in half. Using more slabs can result in even further reductions in the total amount of work. The image shown in Figure 4 was deblurred using two slabs. In one-dimensional experiments, there was very little degradation in the quality of the deblurring when using two slabs, as compared with a single slab.

Acknowledgements

This research was supported by Office of Naval Research Grant N00014-85-K-0077, Work Unit NR 4007006, and NSF grant DCR-8403300. We thank Robert Kohn for providing many of the ideas, and for his helpful guidance.

References

- [1] Poggio, T. and V. Torre, "Ill-posed problems and regularization analysis in early vision," MIT A.I. Laboratory Technical Memo 773 (1984).
- [2] Marr, D., *Vision*, W. H. Freeman and Company (1982).
- [3] Zeevi, Y. Y. and D. Rotem, "Image reconstruction from zero crossings," *IEEE ASSP* 34, pp. 1269-1277 (1986).
- [4] Hummel, Robert, "Representations based on zero-crossings in scale-space," *Proceedings of the IEEE Computer Vision and Pattern Recognition Conference*, pp. 204-209 (June, 1986).

¹Involving a combination of defocusing and vaseline.



Figure 5. An original 64 by 64 image (upper left), and an optically blurred version (upper right). The second row shows an intermediate level of deblurring (left), and the base (deblurred) level (right) after deblurring by equation error minimization.

- [5] John, F., "Numerical solution of the equation of heat conduction at preceeding times," *Annali de Mathematica* 17, pp. 129-142 (1955).
- [6] Shvaytser, H. and S. Peleg, "Pictures as elements in a vector space," *IEEE Conference on Computer Vision and Pattern Recognition*, (June, 1983). See also Shvaytser and Peleg, "Inversion of picture Operators," *IEEE Conference on Systems, Man, and Cybernetics*, Bombay and New Delhi, December 1983, p. 202-206. To appear in *Pattern Recognition Letters*.
- [7] Carasso, A., J. Sanderson, and J. Hyman, "Digital removal of random media image degradations by solving the diffusion equation backwards in time," *SIAM Journal of*

Numerical Analysis 15, pp. 344-367 (1978).

- [8] Hummel, Robert A. and B. Kimia, "Deblurring gaussian blur," *Computer Vision, Graphics, and Image Processing*, (1986). In galleys.
- [9] Courant, R. and D. Hilbert, *Methods of Mathematical Physics*, Interscience Publishers, London (1962).
- [10] Poggio, T., V. Torre, and C. Koch, "Computational vision and regularization theory," *Nature* 317, pp. 314-319 (1985).
- [11] Lowe, Bruce, *A variational method for parameter identification*, PhD Thesis, New York University October, 1986.
- [12] Strang, G. and G. Fix, *An Analysis of the Finite Element Method*, Prentice Hall (1973).
- [13] Ciarlet, Ph.G., in *Numerical Analysis of the Finite Element Method*, Les Presses de l'universite de Montreal (1975).
- [14] Golub, G.H. and C.F. Van Loan, in *Matrix Computations*, John Hopkins Press, Baltimore (1984).

NYU COMPSCI TR-267 c.1
Hummel, Robert A
Solving ill-conditioned
problems by minimizing
equation error.

NYU COMPSCI TR-267 c.1
Hummel, Robert A
Solving ill-conditioned
problems by minimizing
equation error.

APR. 06 1987

DATE DUE

GAYLORD

PRINTED IN U.S.A.

